

Linear Time Lempel-Ziv Factorization: Simple, Fast, Small

Juha Kärkkäinen, Dominik Kempa, and Simon J. Puglisi

Department of Computer Science,
University of Helsinki
Helsinki, Finland
`firstname.lastname@cs.helsinki.fi`

Abstract. Computing the LZ factorization (or LZ77 parsing) of a string is a computational bottleneck in many diverse applications, including data compression, text indexing, and pattern discovery. We describe new linear time LZ factorization algorithms, some of which require only $2n \log n + O(\log n)$ bits of working space to factorize a string of length n . These are the most space efficient linear time algorithms to date, using $n \log n$ bits less space than any previous linear time algorithm. The algorithms are also practical, simple to implement, and very fast in practice.

1 Introduction

In the 35 years since its discovery the LZ77 factorization of a string — named after its authors Abraham Lempel and Jacob Ziv, and the year 1977 in which it was published — has been applied all over computer science. The first uses of LZ77 were in data compression, and to this day it lies at the heart of efficient and widely used file compressors, like `gzip` and `7zip`. LZ77 is also important as a *measure* of compressibility. For example, its size is a lower bound on the size of the smallest context-free grammar that represents a string [1].

In all these applications (and most of the many others we have not listed) computation of the factorization is a time- and space-bottleneck in practice. Our particular motivation is the construction of compressed full-text indexes [13], several recent and powerful instances of which are based on LZ77 [6,5,12].

Related work. There exists a variety of worstcase linear time algorithms to compute the LZ factorization [2,3,4,7,14]. All of them require at least $3n \log n$ bits of working space¹ in the worstcase. The most space efficient linear time algorithm is due to Chen et al. [2]. By overwriting the suffix array it achieves a working space of $(2n + s) \log n$ bits, where s is the maximal size of the stack used in the algorithm. However, in the worstcase $s = \Theta(n)$. Another space efficient solution requiring $(2n + \sqrt{n}) \log n$ bits of space in the worstcase is from [4] but it computes only the lengths of LZ77 factorization phrases. It can be extended to compute the full parsing at the cost of extra $n \log n$ bits.

All of these algorithms rely on the suffix array, which can be constructed in $O(n)$ time and using $(1 + \epsilon)n \log n$ bits of space (in addition to the input string but including the output of size $n \log n$ bits) [9]. This raises the question of whether the space complexity of linear time LZ77 factorization can be reduced from $3n \log n$ bits. In this paper, we answer the question in the affirmative by describing a linear time algorithm using $2n \log n$ bits.

In terms of practical performance, the fastest linear time LZ factorization algorithms are the very recent ones by Goto and Bannai [7], all using at least $3n \log n$ bits of working space. Other

¹ The working space excludes the input string, the output factorization, and $O(\log n)$ terms.

candidates for the fastest algorithms are described by Kempa and Puglisi [10]. Due to nearly simultaneous publication, no comparison between them exists so far. Experiments in this paper put the algorithms of Kempa and Puglisi slightly ahead. Their algorithms are also very space efficient; one of them uses $2n \log n + n$ bits of working space and others even less. However, their worstcase time complexity is $\Theta(n \log \sigma)$ for an alphabet of size σ . More details about these algorithms are given in Section 2.

Our contribution. We describe two linear time algorithms for LZ factorization. The first algorithm uses $3n \log n$ bits of working space and can be seen as a reorganization of an algorithm by Goto and Bannai [7]. However, this reorganization makes it smaller and faster. In our experiments, this is the fastest of all algorithms when the input is not highly repetitive.

The second algorithm employs a novel combinatorial technique to reduce the working space to $2n \log n$ bits, which is at least $n \log n$ bits less than any previous linear time algorithm uses in the worstcase. The space reduction does not come at a great cost in performance. The algorithm is the fastest on some inputs and not far behind the fastest on others.

Both algorithms share several nice features. They are alphabet-independent, using only character comparisons to access the input. They make just one sequential pass over the suffix array, enabling streaming from disk, which would reduce the working space by a further $n \log n$ bits. They are also very simple and easy to implement.

2 Preliminaries

Strings. Throughout we consider a string $X = X[1..n] = X[1]X[2] \dots X[n]$ of $|X| = n$ symbols drawn from an ordered alphabet of size σ .

For $i = 1, \dots, n$ we write $X[i..n]$ to denote the *suffix* of X of length $n - i + 1$, that is $X[i..n] = X[i]X[i+1] \dots X[n]$. We will often refer to suffix $X[i..n]$ simply as “suffix i ”. Similarly, we write $X[1..i]$ to denote the *prefix* of X of length i . We write $X[i..j]$ to represent the *substring* $X[i]X[i+1] \dots X[j]$ of X that starts at position i and ends at position j . Let $\text{lcp}(i, j)$ denote the length of the longest-common-prefix of suffix i and suffix j . For example, in the string $X = \text{zzzzzipzip}$, $\text{lcp}(2, 5) = 1 = |z|$, and $\text{lcp}(5, 8) = 3 = |\text{zip}|$. For technical reasons we define $\text{lcp}(i, 0) = \text{lcp}(0, i) = 0$ for all i .

Suffix Arrays. The suffix array SA is an array $\text{SA}[1..n]$ containing a permutation of the integers $1..n$ such that $X[\text{SA}[1]..n] < X[\text{SA}[2]..n] < \dots < X[\text{SA}[n]..n]$. In other words, $\text{SA}[j] = i$ iff $X[i..n]$ is the j^{th} suffix of X in ascending lexicographical order. The inverse suffix array ISA is the inverse permutation of SA , that is $\text{ISA}[i] = j$ iff $\text{SA}[j] = i$. Conceptually, $\text{ISA}[i]$ tells us the position of suffix i in SA .

The array $\Phi[0..n]$ (see [8]) is defined by $\Phi[i] = \text{SA}[\text{ISA}[i] - 1]$, that is, the suffix $\Phi[i]$ is the immediate lexicographical predecessor of the suffix i . For completeness and for technical reasons we define $\Phi[\text{SA}[1]] = 0$ and $\Phi[0] = \text{SA}[n]$ so that Φ forms a permutation with one cycle.

LZ77. The LZ77 factorization uses the notion of a *longest previous factor* (LPF). The LPF at position i in X is a pair (p_i, ℓ_i) such that, $p_i < i$, $X[p_i..p_i + \ell_i - 1] = X[i..i + \ell_i - 1]$ and $\ell_i > 0$ is maximized. In other words, $X[i..i + \ell_i - 1]$ is the longest prefix of $X[i..n]$ which also occurs at some position $p_i < i$ in X . If $X[i]$ is the leftmost occurrence of a symbol in X then such a pair does not exist. In this case we define $p_i = X[i]$ and $\ell_i = 0$. Note that there may be more than one potential p_i , and we do not care which one is used.

The LZ77 factorization (or LZ77 parsing) of a string X is then just a greedy, left-to-right parsing of X into longest previous factors. More precisely, if the j th LZ factor (or *phrase*) in the parsing is to start at position i , then we output (p_i, ℓ_i) (to represent the j th phrase), and then the $(j + 1)$ th phrase starts at position $i + \ell_i$, unless $\ell_i = 0$, in which case the next phrase starts at position $i + 1$. We call a factor (p_i, ℓ_i) *normal* if it satisfies $\ell_i > 0$ and *special* otherwise. The number of phrases in the factorization is denoted by z .

For the example string $X = \text{zzzzzipzip}$, the LZ77 factorization produces:

$$(z, 0), (1, 4), (i, 0), (p, 0), (5, 3).$$

The second and fifth factors are normal, and the other three are special.

NSV/PSV. The LPF pairs can be computed using *next and previous smaller values* (NSV/PSV) defined as

$$\text{NSV}_{\text{lex}}[i] = \min\{j \in [i + 1..n] \mid \text{SA}[j] < \text{SA}[i]\}$$

$$\text{PSV}_{\text{lex}}[i] = \max\{j \in [1..i - 1] \mid \text{SA}[j] < \text{SA}[i]\}.$$

If the set on the right hand side is empty, we set the value to 0. Further define

$$\text{NSV}_{\text{text}}[i] = \text{SA}[\text{NSV}_{\text{lex}}[\text{ISA}[i]]] \tag{1}$$

$$\text{PSV}_{\text{text}}[i] = \text{SA}[\text{PSV}_{\text{lex}}[\text{ISA}[i]]]. \tag{2}$$

If $\text{NSV}_{\text{lex}}[\text{ISA}[i]] = 0$ ($\text{PSV}_{\text{lex}}[\text{ISA}[i]] = 0$) we set $\text{NSV}_{\text{text}}[i] = 0$ ($\text{PSV}_{\text{text}}[i] = 0$).

If (p_i, ℓ_i) is a normal factor, then either $p_i = \text{NSV}_{\text{text}}[i]$ or $p_i = \text{PSV}_{\text{text}}[i]$ is always a valid choice for p_i [3]. To choose between the two (and to compute the ℓ_i component), we have to compute $\text{lcp}(i, \text{NSV}_{\text{text}}[i])$ and $\text{lcp}(i, \text{PSV}_{\text{text}}[i])$ and choose the larger of the two, see Fig. 1.

Algorithm LZ-Factor(i, psv, nsu)

```

1: if  $\text{lcp}(i, psv) > \text{lcp}(i, nsu)$  then
2:    $(p, \ell) \leftarrow (psv, \text{lcp}(i, psv))$ 
3: else
4:    $(p, \ell) \leftarrow (nsu, \text{lcp}(i, nsu))$ 
5: if  $\ell = 0$  then  $p = X[i]$ 
6: output factor  $(p, \ell)$ 
7: return  $i + \max(\ell, 1)$ 
```

Fig. 1. The basic procedure for computing a phrase starting at a position i given $psv = \text{PSV}_{\text{text}}[i]$ and $nsu = \text{NSV}_{\text{text}}[i]$. The return value is the starting position of the next phrase.

Lazy LZ Factorization. The fastest LZ factorization algorithms in practice are from recent papers by Kempa and Puglisi [10] and Goto and Bannai [7]. A common feature between them is a lazy evaluation of LCP values: $\text{lcp}(i, \text{NSV}_{\text{text}}[i])$ and $\text{lcp}(i, \text{PSV}_{\text{text}}[i])$ are computed only when i is a starting position of a phrase. The values are computed by a plain character-by-character comparison of the suffixes, but it is easy to see that the total time complexity is $O(n)$. This is in contrast to most previous algorithms that compute the LCP values for every suffix using more complicated techniques. The new algorithms in this paper use lazy evaluation too.

Goto and Bannai [7] describe algorithms that compute and store the full set of NSV/PSV values. One of their algorithms, BGT, computes the NSV_{text} and PSV_{text} arrays with the help of the Φ array. The LZ factorization is then easily computed by repeatedly calling **LZ-Factor**. Two other algorithms, BGS and BGL, compute the NSV_{lex} and PSV_{lex} arrays and use them together with SA and ISA to simulate NSV_{text} and PSV_{text} as in Eqs. (1) and (2). All three algorithms run in linear time and they use $3n \log n$ (BGT), $4n \log n$ (BGL) and $(4n + s) \log n$ (BGS) bits of working space, where s is the size of the stack used by BGS. In the worst case $s = \Theta(n)$. The algorithms for computing the NSV/PSV values are not new but come from [14] (BGT) and from [3] (BGL and BGS). However, the use of lazy LCP evaluation makes the algorithms of Goto and Bannai faster in practice than earlier algorithms.

Kempa and Puglisi [10] extend the lazy evaluation to the NSV/PSV values too. Using ISA and a small data structure that allows arbitrary NSV/PSV queries over SA to be answered quickly, they compute $\text{NSV}_{\text{text}}[i]$ and $\text{PSV}_{\text{text}}[i]$ only when i is a starting position of a phrase. The approach requires $(2 + 1/b)n \log n$ bits of working space and $O(n + zb + z \log(n/b))$ time, where b is a parameter controlling a space-time tradeoff in the NSV/PSV data structure. If we set $b = \log n$, and given $z = O(n / \log_{\sigma} n)$, then in the worstcase the algorithm requires $O(n \log \sigma)$ time, and $2n \log n + n$ bits of space. Despite the superlinear time complexity, this algorithm (ISA9) is both faster and more space efficient than earlier linear time algorithms. Kempa and Puglisi also show how to reduce the space to $(1 + \epsilon)n \log n + n + O(\sigma \log n)$ bits by storing a succinct representation of ISA (algorithms ISA6r and ISA6s). Because of the lazy evaluation, these algorithms are especially fast when the resulting LZ factorization is small.

3 $3n \log n$ -Bit Algorithm

Our first algorithm is closely related to the algorithms of Goto and Bannai [7], particularly BGT and BGS. It first computes the PSV_{text} and NSV_{text} arrays and uses them for lazy LZ factorization similarly to the BGT algorithm. However, the NSV/PSV values are computed using the technique of the BGS algorithm, which comes originally from [3]. The algorithm is given Figure 2.

Algorithm KKP3

```

1: SA[0] ← 0      // bottom of stack
2: SA[n + 1] ← 0  // empties the stack at end
3: top ← 0        // top of stack
4: for i ← 1 to n + 1 do
5:   while SA[top] > SA[i] do
6:     NSVtext[SA[top]] ← SA[i]
7:     PSVtext[SA[top]] ← SA[top - 1]
8:     top ← top - 1    // pop from stack
9:   top ← top + 1
10:  SA[top] ← SA[i]    // push to stack
11: i ← 1
12: while i ≤ n do
13:   i ← LZ-Factor(i, PSVtext[i], NSVtext[i])

```

Fig. 2. LZ factorization using $3n \log n$ bits of working space (the arrays SA, NSV_{text} and PSV_{text}).

The advantages of our algorithm compared to those of Goto and Bannai are:

1. All of the algorithms of Goto and Bannai use an auxiliary array of size n , either ISA or Φ . We need no such auxiliary array, which saves both space and time.
2. Both BGS and our algorithm need a stack whose maximum size is not known in advance and can be $\Theta(n)$ in the worst case. BGS uses a dynamically growing separate stack while we overwrite the suffix array with the stack. This is possible because our algorithm makes just one pass over the suffix array (like BGT but unlike BGS) and the stack is never larger than the already scanned part of SA .
3. Similar to the algorithms of Goto and Bannai, we store the arrays PSV_{text} and NSV_{text} interleaved so that the values $\text{PSV}_{\text{text}}[i]$ and $\text{NSV}_{\text{text}}[i]$ are next to each other. We compute the PSV value when popping from the stack instead of when pushing to the stack as BGS does. This way $\text{PSV}_{\text{text}}[i]$ and $\text{NSV}_{\text{text}}[i]$ are computed and written at the same time which can reduce the number of cache misses.

4 $2n \log n$ -Bit Algorithm

Our second algorithm reduces space by computing and storing only the NSV values at first. It then computes the PSV values from the NSV values on the fly. As a side effect, the algorithm also computes the Φ array!

For $t \in [0..n]$, let $\mathcal{X}_t = \{X[i..n] \mid i \leq t\}$ be the set of suffixes starting at or before position t . Let Φ_t be Φ restricted to \mathcal{X}_t , that is, for $i \in [1..t]$, suffix $\Phi_t[i]$ is the immediate lexicographical predecessor of suffix i among the suffixes in \mathcal{X}_t . In particular, $\Phi_n = \Phi$. As with the full Φ , we make Φ_t a complete unicyclic permutation by setting $\Phi_t[i_{\min}] = 0$ and $\Phi_t[0] = i_{\max}$, where i_{\min} and i_{\max} are the lexicographically smallest and largest suffixes in \mathcal{X}_t . We also set $\Phi_0[0] = 0$. A useful way to view Φ_t is as a circular linked list storing \mathcal{X}_t in the descending lexicographical order with $\Phi_t[0]$ as the head of the list.

Now consider computing Φ_t given Φ_{t-1} . We need to insert a new suffix t into the list, which can be done using standard insertion into a singly-linked list provided we know the position. It is easy to see that t should be inserted between $\text{NSV}_{\text{text}}[t]$ and $\text{PSV}_{\text{text}}[t]$. Thus

$$\Phi_t[i] = \begin{cases} t & \text{if } i = \text{NSV}_{\text{text}}[t] \\ \text{PSV}_{\text{text}}[t] & \text{if } i = t \\ \Phi_{t-1}[i] & \text{otherwise} \end{cases}$$

and furthermore

$$\text{PSV}_{\text{text}}[t] = \Phi_{t-1}[\text{NSV}_{\text{text}}[t]] .$$

The pseudocode for the algorithm is given in Figure 3. The NSV values are computed essentially the same way as in the first algorithm (lines 1–9) and stored in the array Φ . In the second phase, the algorithm maintains the invariant that after t rounds of the loop on lines 12–18, $\Phi[0..t] = \Phi_t$ and $\Phi[t+1..n] = \text{NSV}_{\text{text}}[t+1..n]$.

5 Getting Rid of the Stack

The above algorithms overwrite the suffix array with the stack, which can be undesirable. First, we might need the suffix array later for another purpose. Second, since the algorithms make just one

Algorithm KKP2s

```

1:  $SA[0] \leftarrow 0$  // bottom of stack
2:  $SA[n+1] \leftarrow 0$  // empties the stack at end
3:  $top \leftarrow 0$  // top of stack
4: for  $i \leftarrow 1$  to  $n+1$  do
5:   while  $SA[top] > SA[i]$  do
6:      $\Phi[SA[top]] \leftarrow SA[i]$  //  $\Phi[SA[top]] = NSV_{\text{text}}[SA[top]]$ 
7:      $top \leftarrow top - 1$  // pop from stack
8:    $top \leftarrow top + 1$ 
9:    $SA[top] \leftarrow SA[i]$  // push to stack
10:  $\Phi[0] \leftarrow 0$ 
11:  $next \leftarrow 1$ 
12: for  $t \leftarrow 1$  to  $n$  do
13:    $nsv \leftarrow \Phi[t]$ 
14:    $psv \leftarrow \Phi[nsv]$ 
15:   if  $t = next$  then
16:      $next \leftarrow \text{LZ-Factor}(t, psv, nsv)$ 
17:    $\Phi[t] \leftarrow psv$ 
18:    $\Phi[nsv] \leftarrow t$ 

```

Fig. 3. LZ factorization using $2n \log n$ bits of working space (the arrays SA and Φ).

sequential pass over the suffix array, we could stream the suffix array from disk to further reduce the memory usage. In this section, we describe variants of our algorithms that do not overwrite SA (and still make just one pass over it).

The idea is to replace the stack with PSV_{text} pointers. If j is the suffix on the top of the stack, then the next suffixes in the stack are $PSV_{\text{text}}[j]$, $PSV_{\text{text}}[PSV_{\text{text}}[j]]$, etcetera. This can be easily seen in how the PSV_{text} values are computed in KKP3 (line 7 in Fig. 2). Thus given PSV_{text} we do not need an explicit stack at all. Both of our algorithms can be modified to exploit this:

- In KKP3, we need to compute the PSV_{text} values when pushing on the stack rather than when popping. The body of the main loop (lines 5–10 in Fig. 2) now becomes:

```

while  $top > SA[i]$  do
   $NSV_{\text{text}}[top] \leftarrow SA[i]$ 
   $top \leftarrow PSV_{\text{text}}[top]$ 
 $PSV_{\text{text}}[SA[i]] \leftarrow top$ 
 $top \leftarrow SA[i]$ 

```

- KKP2s needs to be modified to compute PSV_{text} values first instead of NSV_{text} values. The PSV_{text} -first version is symmetric to the NSV_{text} -first algorithm. In particular, Φ_t is replaced by the inverse permutation Φ_t^{-1} . The algorithm is shown in Fig. 4.

The versions without an explicit stack are slightly slower because of the non-locality of the pointer accesses. If we need to avoid overwriting SA , a faster alternative would be to use a separate stack. However, the stack can grow as big as n (for example when $X = a^{n-1}b$) which increases the worst case space requirement by $n \log n$ bits.

We can get the best of both alternatives by adding a fixed size stack buffer to the stackless version. The buffer holds the top part of the stack to speed up stack operations. When the buffer gets full, the bottom half of its contents is discarded, and when the buffer gets empty, it is filled

Algorithm KKP2n

```

1:  $top \leftarrow 0$  // top of stack
2: for  $i \leftarrow 1$  to  $n$  do
3:   while  $top > SA[i]$  do
4:      $top \leftarrow \Phi^{-1}[top]$  // pop from stack
5:    $\Phi^{-1}[SA[i]] \leftarrow top$  //  $\Phi^{-1}[SA[i]] = PSV_{\text{text}}[SA[i]]$ 
6:    $top \leftarrow SA[i]$  // push to stack
7:  $\Phi^{-1}[0] \leftarrow 0$ 
8:  $next \leftarrow 1$ 
9: for  $t \leftarrow 1$  to  $n$  do
10:   $psv \leftarrow \Phi^{-1}[t]$ 
11:   $nsv \leftarrow \Phi^{-1}[psv]$ 
12:  if  $t = next$  then
13:     $next \leftarrow \text{LZ-Factor}(t, psv, nsv)$ 
14:     $\Phi^{-1}[t] \leftarrow nsv$ 
15:     $\Phi^{-1}[psv] \leftarrow t$ 

```

Fig. 4. LZ factorization using $2n \log n$ bits of working space (the arrays SA and Φ^{-1}) without an explicit stack. The SA remains intact after the computation.

half way using the PSV pointers. This version is called KKP2b. The time complexity remains linear and is independent of the buffer size.

6 Experimental Results

We implemented the algorithms described in this paper and compared their performance in practice to algorithms from [10] and [7]. Experiments measured the time to compute the LZ factorization of the text. All algorithms take the text and the suffix array as an input hence we omit the time to compute SA . The data set used in experiments is described in detail in Table 1.

Experiments Setup. We performed experiments on a 2.4GHz Intel Core i5 CPU equipped with 3072KB L2 cache and 4GB of main memory. The machine had no other significant CPU tasks running and only a single thread of execution was used. The OS was Linux (Ubuntu 10.04, 64bit) running kernel 2.6.32. All programs were compiled using g++ version 4.4.3 with `-O3 -static -DNDEBUG` options. For each combination of algorithm and test file we report the median runtime from five executions. The times were recorded with the standard C `clock` function. All data structures reside in main memory during computation.

Discussion. In nearly all cases algorithms introduced in this paper outperform the algorithms from [7] (which are, to our knowledge, the fastest up-to-date linear time LZ factorization algorithms) while using the same or less space. In particular the KKP2 algorithms are always faster and simultaneously use at least $n \log n$ bits less space. A notably big difference is observed for non-repetitive data, where KPP3 significantly dominates all prior solutions.

The new algorithms (e.g. KKP2b) also dominate in most cases the general purpose practical algorithms from [10] (ISA9 and ISA6s), while offering stronger worst case time guarantees, but are a frame slower (and use about 50% more space in practice) than ISA6r for highly repetitive data.

The comparison of KKP2n to KKP2s reveals the expected slowdown (up to 16%) due to the non-local stack simulation. However, this effect is almost completely eliminated by buffering the

Name	$n/2^{20}$	σ	n/z	Source	Description
proteins	150	25	9.57	S	Swissprot database
english	150	220	13.77	S	Gutenberg Project
dna	150	16	14.65	S	Human genome
sources	150	228	17.67	S	Linux and GCC sources
coreutils	150	236	110	R/R	$9 \times$ GNU Coreutils source
cere	150	5	112	R/R	$37 \times$ baking yeast genome
kernel	150	160	214	R/R	$36 \times$ Linux Kernel sources
einstein.en	150	124	3634	R/R	Wikipedia
proteins.001.1	100	21	295	R/PR	$100 \times$ 1MB protein
english.001.2	100	106	312	R/PR	$100 \times$ 1MB english
dna.001.1	100	5	340	R/PR	$100 \times$ 1MB dna
sources.001.2	100	98	355	R/PR	$100 \times$ 1MB sources
tm29	150	2	2912K	R/A	Thue-Morse sequence
rs.13	150	2	3024K	R/A	Run-Rich String sequence

Table 1. Files used in the experiments. The files are from the standard (S) Pizza&Chili corpus (<http://pizzachili.dcc.uchile.cl/texts.html>) and from the repetitive (R) Pizza&Chili corpus (<http://pizzachili.dcc.uchile.cl/rep corpus.html>). The repetitive corpus consists of files containing multiple copies of similar data (R), artificially generated sequences (A), and files created from standard corpus by concatenating 100 copies of 1MB prefix and mutating them randomly (PR). The value of n/z (average length of phrase in LZ factorization) is included as measure of repetitiveness.

top part of the stack (KKP2b). With a 256KB buffer we obtained runtimes almost identical to KKP2s ($< 1\%$ difference in all cases). We observed a similar effect when applying this optimization to the KKP3 algorithm but, for brevity, we only present the improvement for KKP2.

Finally, we observe that KKP2b, despite being slower than KPP3 on non-repetitive data, has runtimes that are very close to best in each category, making it perhaps the most applicable general purpose algorithm.

Testfile	KKP3	KKP2s	KKP2b	KKP2n	ISA6r	ISA6s	ISA9	iBGS	iBGL	iBGT
proteins	74.5	83.9	84.1	88.1	-	198.0	92.7	99.8	123.2	171.4
english	75.7	80.6	80.6	84.6	-	171.0	83.9	93.2	108.6	153.9
dna	81.7	92.7	92.7	97.3	-	175.2	86.1	97.5	113.4	188.0
sources	50.5	54.7	54.8	56.1	-	115.0	59.3	69.3	77.8	99.8
coreutils	43.6	40.2	40.2	40.6	43.3	49.4	41.9	51.5	52.2	55.4
cere	63.2	53.3	53.2	57.7	51.8	56.3	53.0	65.5	66.1	84.1
kernel	45.7	41.6	41.5	42.2	39.2	45.7	42.8	52.9	53.0	56.2
einstein.en	56.9	43.6	43.5	47.6	31.1	37.1	45.2	60.0	58.6	52.8
proteins.001.1	52.6	43.1	43.1	50.0	40.7	45.3	46.6	58.4	57.6	59.6
english.001.2	52.0	43.4	43.8	52.2	40.4	45.1	45.3	57.7	56.0	79.4
dna.001.1	55.6	43.9	43.9	50.8	39.2	43.7	45.0	58.5	57.8	62.5
sources.001.2	43.1	40.5	40.5	47.8	38.0	42.8	41.1	54.2	52.4	72.2
tm29	38.2	35.1	35.1	38.7	34.2	39.6	36.4	44.1	44.2	44.4
rs.13	77.8	49.0	49.4	52.0	34.8	40.8	51.8	59.5	59.5	56.5

Table 2. Times for computing LZ factorization. The times are seconds per gigabyte and do not include any reading from or writing to disk.

7 Future Work

For data of low to medium repetitiveness the algorithms introduced in this paper are the fastest available. These algorithms should adapt easily to a semi-external setting because, apart from the need to permute the NSV/PSV values into text order, which can be handled in external memory, all non-sequential memory accesses are restricted to the input string. We are currently exploring this direction.

There are several interesting open problems. One is the need for a fully external memory algorithm for LZ factorization, especially given the recent pattern matching indexes which use LZ77. Relatedly, parallel and distributed approaches are also of high interest. A recent step in the external memory direction is [11].

Another problem is to find a scalable way to accurately estimate the size of the LZ factorization in lieu of actually computing it. Such a tool would be useful for entropy estimation, and to guide the selection of appropriate compressors and compressed indexes when managing massive data sets.

Finally, one wonders if only $(1 + \epsilon)n \log n + O(\log n)$ bits of working memory is enough for linear runtime. The most space-efficient algorithm in this paper uses $2n \log n + O(\log n)$ bits, and in [10] working space of $(1 + \epsilon)n \log n + n + O(\sigma \log n)$ bits (for arbitrary constant ϵ) is achieved, but at the price of $O(n \log \sigma)$ runtime.

Acknowledgments

Our thanks go to Keisuke Goto and Hideo Bannai for sending us a preliminary copy of their manuscript.

References

1. M. Charikar, E. Lehman, D. Liu, R. Panigrahy, M. Prabhakaran, A. Sahai, and a. shelat. The smallest grammar problem. *IEEE Transactions on Information Theory*, 51(7):2554–2576, 2005.
2. G. Chen, S. J. Puglisi, and W. F. Smyth. Fast and practical algorithms for computing all the runs in a string. In *Proc. Symposium on Combinatorial Pattern Matching (CPM)*, pages 307–315, 2007.
3. M. Crochemore and L. Ilie. Computing longest previous factor in linear time and applications. *Information Processing Letters*, 106(2):75–80, 2008.
4. M. Crochemore, L. Ilie, and W. F. Smyth. A simple algorithm for computing the Lempel Ziv factorization. In *Proc. Data Compression Conference (DCC)*, pages 482–488, 2008.
5. T. Gagie, P. Gawrychowski, J. Kärkkäinen, Y. Nekrich, and S. J. Puglisi. A faster grammar-based self-index. In *Proc. Conference on Language and Automata Theory and Applications (LATA)*, LNCS 7183, pages 240–251, 2012.
6. T. Gagie, P. Gawrychowski, and S. J. Puglisi. Faster approximate pattern matching in compressed repetitive texts. In *Proc. Symposium on Algorithms and Computation (ISAAC)*, pages 653–662, 2011.
7. K. Goto and H. Bannai. Simpler and faster Lempel Ziv factorization. <http://arxiv.org/abs/1211.3642>, 2012.
8. J. Kärkkäinen, G. Manzini, and S. J. Puglisi. Permuted longest-common-prefix array. In *Proc. Symposium on Combinatorial Pattern Matching (CPM)*, LNCS 5577, pages 181–192. Springer, 2009.
9. J. Kärkkäinen, P. Sanders, and S. Burkhardt. Linear work suffix array construction. *Journal of the ACM*, 53(6):918–936, 2006.
10. D. Kempa and S. J. Puglisi. Lempel-Ziv factorization: simple, fast, practical. In *Proc. Algorithm Engineering and Experiments (ALENEX)*, 2013. In press.
11. D. Kempa and S. J. Puglisi. Lightweight LZ77 factorization, 2013. Manuscript.
12. S. Kreft and G. Navarro. Self-indexing based on LZ77. In *Proc. Symposium on Combinatorial Pattern Matching (CPM)*, LNCS 6661, pages 41–54, 2011.
13. G. Navarro and V. Mäkinen. Compressed full-text indexes. *ACM Computing Surveys*, 39(1):article 2, 2007.
14. E. Ohlebusch and S. Gog. Lempel-Ziv factorization revisited. In *Proc. Symposium on Combinatorial Pattern Matching (CPM)*, LNCS 6661, pages 15–26, 2011.